DOCUMENT RESUME

ED 431 007 TM 029 839

AUTHOR Habick, Timothy
TITLE Item Variants in Computer-Based Tests.
PUB DATE 1999-04-00
NOTE 12p.; Paper presented at the Annual Meeting of the American Educational Research Association (Montreal, Quebec, Canada, April 19-23, 1999).
PUB TYPE Reports - Descriptive (141) -- Speeches/Meeting Papers (150)
EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS Adaptive Testing; *Computer Assisted Testing; *Item Banks; Test Construction; Test Format; *Test Items

ABSTRACT

With the advent of computer-based testing (CBT) and the need to increase the number of items available in computer adaptive test pools, the idea of item variants was conceived. An item variant can be defined as an item with content based on an existing item to a greater or lesser degree. Item variants were first proposed as a way to enhance test security by increasing the size of CBT item pools. Variants are now also seen as useful rapid item-creation pools in programs that use paper-based testing exclusively because they also enhance security in that situation. This report summarizes the part played by variants in item pool expansion and item pool security, focusing on types of variants and their appropriate uses and the difficulties for item pool management created by item variants. Attachments discuss a logical reasoning item and medium variants and an analytical reasoning stimulus and medium variants. (SLD)

ED 431 007

Item Variants in Computer-Based Tests

Timothy Habick

TM029839

This paper is prepared for the:
Annual Meeting of the American Educational Research Association in Montreal Canada
April 1999

Item Variants in Computer-Based Tests

Timothy Habick
Educational Testing Service
April 23, 1999

Test item creation has traditionally been performed at Educational Testing Service in a manner similar to the way that most professional writers approach their work. Each piece of writing had to be an original creative product that was ideally written from scratch by a single writer, who would make revisions suggested by reviewers and editors but who never relinquished authorship of the original product. This conception of item writing has now changed.

With the advent of computer-based testing (CBT) and the need to increase significantly the number of items available in computer-adaptive test pools, the idea of item variants was conceived. The first such large-scale project was initiated in June, 1994. The purpose of this project was to write a large number of variants based on practice book questions from the Graduate Record Examinations (GRE) General Test, to pretest them, and then to include them into the GRE operational pools.

Especially when combined with the practice of continuous testing, computer-based testing has raised significant test security concerns. Continuous testing is a test-administration method that allows consumers the convenience of scheduling a test for a time of their choosing; in many locations testing is available six days a week. In the past, large paper-based test administrations allowed on a single day the exposure of a relatively small number of test items to a very large number of test takers. For the GRE General Test, for example, a single paper-based administration typically attracted from 50,000 to 100,000 test takers in the United States alone. Although the number of base test forms was limited, test security procedures for paper-based administrations ensured that not all test takers were presented with identical test booklets. In addition, test takers appearing for the next paper-based test administration of the GRE would be unlikely to see any of the questions used in the preceding administration. Test forms might be used a second or third time (often overseas or in some other testing environment) in some random sequence that was difficult for a test taker or test-preparation school to predict. In some cases, the test form might be disclosed immediately after its first operational use.

In a continuous-testing CBT environment, however, a single large set of operational questions could be reused for several days at the same location, although any given subset would be seen by only a small number of test takers. Without additional controls, the security of those test items could be diminished even after the first day of CBT exposure. Item variants are one aspect of the additional controls used to ensure the security of CBT pools.

An item variant can be defined as an item whose content is based on an existing item to a greater or lesser degree. The variant may be expected to have significantly distinct psychometric properties or may be expected to perform almost identically to the base

item. Item creation specialists produce different types of variants and variant families in response to the changing item pool needs of a particular program.

Item variants were first proposed as a way to enhance security by increasing the size of CBT item pools. Concerns were soon raised, however, that exposure of a given item might threaten the security of members of the variant family. However, far from decreasing the security of CBT pools, item variants were soon seen as increasing security by making it obviously counterproductive for a test taker to use hearsay information about the content of items. Some members of a variant family, for example, might be so close at first glance but so different in their deeper logical structure that test takers who saw a given item in an operational test and then tried to explain the item and its answer to an acquaintance would actually mislead the acquaintance, who would be more likely to encounter a variant of the item than the base item itself. The GRE program has taken advantage of the security benefits of variants by prominently announcing in its information bulletin and on its web site the fact that modified versions of questions may appear in any testing session. Test takers are likewise informed that modified versions of items included in official GRE publications may also be included in a test they take. Test takers are advised to attend carefully to the text of each item as it is presented and not to attempt to answer the question on the basis of some memory of a similar item that they might have seen or heard about.

This report will summarize the part played by variants in item pool expansion and item pool security.

Item variants as an innovation

From one point of view, there is nothing new about item variants. For certain types of items, especially pure mathematics items, item variants have always been created and incorporated judiciously into different test editions. Items testing a given aspect of algebra, for example, might be expressed in terms of variables x and y or variables p and r. A divisor might be 2 or 3. However, an item testing the skill at issue would be essentially the same, and expressed in essentially the same way, as any other item. Similarly, in an item type like logical reasoning, sometimes more than one question was asked about a single, invariant stimulus argument. The multiple items attached to a given stimulus were, in fact, item variants--at least in the sense that they exploited the potential of the stimulus to generate more than one item.

From another point of view, item variants have revolutionized item creation at ETS, especially for the most expensive item types. The qualities of item uniqueness and single authorship that characterized item production only a few years ago are seen today not only as wasteful of resources but also as quite inappropriate for the task at hand. Test takers, of course, do not have access to the entire item pool (or "vat" from which individual computer-adaptive pools are formed) but typically encounter only one member of an item variant family. (Indeed, tracking of variant family members' identification numbers is employed in order to help prevent a situation in which a single test taker is presented with more than one item variant family member in a given operational test. Test takers might, however, see an item in an official ETS practice test book or software

package and then encounter a variant of that item operationally.) Therefore, test takers do not experience the somewhat numbing effect of seeing several questions on the same topic asked in slightly different or very different ways. On the surface, the test taker encounters a test that looks quite similar to a test composed of unique items from unique sources. Beneath this exterior, however, lies a collection of slightly distinct and maximally distinct items with a variety of psychometric properties. These variant-family members are not intended to be seen (and never are seen) by any one test taker. Test takers are presented with only those items whose psychometric properties are most appropriate for them, according to their performance on the items already answered in the particular section of a given computer-adaptive testing session.

Types of item variants

Item variants can be classified in different ways. For example, it is useful to classify variants as close, medium, and far variants. Close variants are easily recognizable as belonging to the same variant family by both test developers and test takers. Medium variants are recognizable as belonging to the same variant family by test developers and may also be recognizable in this way to test takers. However, medium variants are expected to have psychometric properties that are distinct from their family members, and many test takers might not immediately notice the variant-family relationship. Far variants are items that might have been based on an original model but have undergone such extensive changes that even test developers might not notice the variant-family relationship. Far variants are therefore not treated as variants at all but rather as independent items. *Far variants are only variants in terms of their writing history, and are not to be understood as included when the word "variants" is used below.*

It can also be useful (especially for mathematics items) to classify variants as generic or nongeneric. Generic items and variants assess a specific skill in a standard or iconic way. That is, any item that attempts to measure a certain specific skill will by its very nature have a certain recognizable appearance: only the variables used in the item will be different. Generic items are inherently unaffected by certain types of security issues. The relevant issue for test takers is not what would a question assessing a certain skill look like but whether a question of this genre will appear in the particular form of the test they happen to take. Therefore, generic items are said, somewhat paradoxically, to have low memorability precisely because they are so common. It is easy to remember that there is a certain type of equation in the examination, but hard to remember exactly which variables and values were used. Hearsay information from one test taker to a future test taker that a certain type of generic item appeared in the test, therefore, would be no more valuable to the future test taker than a review of similar generic items available in disclosed tests that are easily available to the public. Nongeneric items, on the other hand, are items that are memorable because their content is in some way unique. Typically, a nongeneric item involves a story line that is easily told and retold. For example, hearsay information about a nongeneric item might be expressed in this way: "There is a question in the test about the dwindling numbers of turtles that nest on Paco Island, and the answer is that feral dogs introduced by recent settlers wait for the turtles to bury their eggs on the beach and then the dogs dig up the nests." There is a clear security advantage to producing variants of nongeneric items; there is some security advantage to producing variants of

generic items, but all items of this type are close variants in any case. For generic items, a testing program simply needs to decide how many items of a particular genre it desires in its repository.

Suitability of particular item types for item-variant creation

Certain types of items are more amenable to variant creation than are other types. Also, it is more cost effective to make certain types of items into variants than it is for other types. As testing programs have included variant creation in their work plans, different procedures have been established for particular item types.

Reading comprehension sets of questions, for example, are not seen as appropriate for variant creation. From one point of view, the different questions attached to a reading comprehension stimulus are themselves already members of a variant family. The difference in this case is that test takers are expected to be exposed to several such reading-set members, and each member is typically considered independent in content from all other set members. The major reason, however, for avoiding variant creation for reading comprehension sets is that making a variant of the reading passage itself would be an extremely time-consuming task, if it could be done reasonably well at all. Limitations on the security benefit of increasing the numbers of items attached to a given reading passage also limit the number of items that can be created from a given passage. Moreover, for reading comprehension sets, exposure of the stimulus itself is a more serious breach of security than is exposure of any particular item. Thus, the expansion of the numbers of items attached to a given passage is severely limited, for practical reasons.

In general, the costlier an item type is, the better the cost-benefit ratio of writing variants for items of that type--and vice versa. For this reason, variants are not created for some relatively inexpensive item types, such as those that assess knowledge of basic grammar for students of English as a second language. As with some basic mathematics items, such items are already generic variants--the same grammatical rule or morphological structure is tested, but each time with a different sample sentence. For other relatively inexpensive item types, such as analogies or antonyms, however, the program's repository of items can be usefully and quickly expanded by variant creation. A potential test taker who hears that the word *taciturn* is included in an item pool with its antonym in a given item listed as *verbose* might not be advantaged when encountering a variant of the item in an operational pool when the antonym key is *garrulous*. Large numbers of such discrete verbal item types can be and have been added to the item repositories.

Perhaps the most useful contribution of variants to bringing down the cost of item writing has occurred with the logical reasoning and analytical reasoning item types. (See examples of these item types attached at the end of this paper.) For logical reasoning, the advantage is clear, since the major cost in item creation is in the identification of an appropriate source topic--typically one that contains or implies an argument. Once a stimulus argument has been identified and written, slight variations of the fact structure woven into the argument could lead to different correct answers, and to psychometrically and impressionistically distinct items. Two major types of variants can be created from a given source. In some variants, the stimulus can remain identical to that of the source

item; only the question asked about the stimulus would be different. For example, one item might ask the test taker to identify a statement that, if true, would most strengthen the argument, and another item might ask for a statement that would most weaken the force of the argument. Another type of variant would involve a different positioning or expression of the fact structure of the argument itself. In this case, the argument, though impressionistically still about the same topic as the original, is structured such that the statement that weakens the variant argument would be different from the one that weakened the source argument. In many cases, a given statement could be included as a possible answer choice in both items; in one, it would be the key, and in another, it would be a distracter.

For the analytical reasoning item type, variant creation has also served well to expand item pools. Both close and medium variants are regularly created. Close variants have logical structures identical to the source sets but use different stories or "clothes" to make the structures appear in the guise of real-world phenomena. For example, a source set might involve a manager assigning workers to particular offices according to a set of rules (If Gary is assigned to Room 201, Fay must be assigned to Room 202, etc.). In a close-variant set, a camp director might be assigning students to ride in different colored boats, according to a set of logically identical rules (If Harry is assigned to the blue boat, Gilda must be assigned to the red boat, etc.). Once a close-variant source set has been created, the variant sets can be written at only minimal additional expense.

Medium variants in analytical reasoning also are created. These types of variant sets take a source set and alter its logical structure only slightly. This ensures that the set will have quite distinct psychometric properties and also makes it very unlikely that a test taker's exposure to the source set or one of its members would significantly affect the individual's behavior when encountering a variant set in an operational setting. The effect of this exposure, in any case, would probably not be perceptibly higher than would the effect of exposure to one of the many similar analytical reasoning sets that have been disclosed to the public in practice books.

For analytical reasoning, the particular mixture in the operational repository of (1) completely distinct sets without variant family members, (2) close variant sets, and (3) medium variant sets has helped to keep the analytical reasoning pool both robust and secure. At the same time, test takers have the same psychometric experience with the item type as they would if all of the sets in the vat were maximally distinct--without any variants at all. Thus, cost savings (over the cost of writing similar numbers of items from scratch) have been realized but without any perceptible negative consequences from a psychometric point of view.

Writing new items with item variant creation in mind

It was soon discovered that the creation of variants would proceed more smoothly (at least for some item types) if original items were written with future variant creation in mind. In this way, test developers ensure that potential base items and sets will not have peculiarities that would make variant creation at a later time problematical or more subject to error. Instead, they create their items or set stimuli with preplanned variable

slots. In some cases with close variants, only the variables need to be changed, by automatic rule (change Gary to Harry; change Gilda to Fay). In other cases, the sentence structures might need to be slightly changed in order to accommodate the idiomatic requirements of a particular verb. (Gary might be assigned to *work in* Room 201; Harry might be assigned to *pilot* the blue boat.)

Writing variants on the basis of items with known statistical properties

Careful consideration of the statistical performance of items in general and item variants in particular led to the realization that it is best to create some types of variants after the parent item has been pretested. This works especially well with analytical reasoning sets. It was realized that there is no advantage to creating a variant family of sets for an analytical reasoning set until the parent set itself has been pretested and has earned good statistics. This is because a close variant of an analytical reasoning set is expected to--and by and large does--perform quite similarly to the parent set. Close variants of sets with good statistical performance will also have similarly good statistical performance, and sets that earned poor statistics cannot normally be expected to produce variant families whose statistics would improve. In fact, research is underway to investigate whether close variants of source sets with proven statistics need to be pretested at all, or if so, whether they need to utilize the same candidate volume as newly written sets.

For logical reasoning items, however, a different situation obtains. Experience has shown that a variant of a logical reasoning item is quite likely to function differently from its source item. The statistical survival of a given item does not necessarily imply the survival of its variant-family members, and vice versa. Therefore, variants of logical reasoning items are created in most cases before the source item is pretested. Test developers work on item-production teams and brainstorm ideas for item variants, attempting to maximize--up to a given limit--the number of items than can be derived from a given source. Writing all logical reasoning variants at the same time as the wording of the source item is being finalized is also efficient because at that time test developers are most conversant with the details of the item's source material. Pretesting all item variants together with their source item helps to ensure that a given test taker will not encounter the source item in an operational portion of the test and one of its variants in an unscored or pretest portion of the test.

Difficulties for pool management created by item variants

Tracking

The existence of item variants in operational test pools does entail some additional costs. Time and care are necessary to ensure that information concerning the variant-family relationship of each item is recorded accurately and is accessible conveniently. Procedures ensuring such efficient use of item pools have been put into place. Still the entry of data and maintenance of that data repository takes time and costs money. This is, in itself, one reason why the use of variants is attractive mainly for the most costly item types. Items that are inexpensive to create might become more costly overall if the logistical demands of variant-family tracking are included in the workload.

Pool diversity
Although helpful in building the absolute numbers of items in pools or the vat, large numbers of variants impose certain natural limitations on the pool--or vat--itself. A pool containing 500 items of which 200 items are members of variant families and 300 items have no variant-family relationship to any other item in the pool is less flexible than a pool containing 500 items of which no item has any variant-family relationship to any other item in the pool. Once the computer-adaptive algorithm has selected a given item for presentation to a test taker, neither that item nor any of its variant-family members can be available for future selection. Therefore, if a given pool contains a high number of variant-family members, the pool itself might not function optimally or at all in the operational setting. This type of negative consequence can be avoided while still preserving the security and cost benefits of variants by not including more than one member of any given variant family in any particular operational pool. In this way, the number of items included in any given pool will be closer to the number of items from which the algorithm can actually select when deciding which item to present next to a test taker.

Disclosure
In compliance with relevant state laws and the practice of the testing program, a certain number of test questions, intact tests, or computer-adaptive test pools are regularly disclosed to the public. The existence of variants in testing pools does complicate the disclosure process. The question is whether the disclosure of a given variant family member would require the disclosure (or retirement) of all other members of the variant family. This depends on the type of variant. Some close variants are so similar to other members of their variant family and their content is so uniquely memorable that disclosure of one family member entails disclosure of the others.

The natural next question to ask is what security or other benefit such close variants might have afforded in the first place. From one point of view, variant families that require or entail disclosure of the entire family when one member is disclosed might not have constituted at the outset a reasonable investment of the program's resources. From another point of view, the existence of close variants in a pool does enhance security in the short term. This short-term value consists in the fact that hearsay information regarding a given variant-family member will be unhelpful to test takers who try to apply that knowledge to another item from that family. When a nongeneric variant-family member is formally disclosed and available for close inspection by future test takers, however, it would be much easier for them to gain an unfair advantage if they encountered a close variant of the disclosed item in an operational test. They would then understand that this variant is just another way of asking the same question.

The long-term value of variant-family members depends, in part, on the memorability and uniqueness of each item's content. In actual practice, it is advisable to disclose or retire some or all close variant-family members in some cases when one member is disclosed, and in other cases disclosure of one item does not entail the disclosure of any variants of it. Effective planning before disclosure decisions are made can prevent the unnecessary disclosure of hundreds of items by choosing, as possible, those items for disclosure that

would not necessitate the disclosure (or retirement) of other items with item parameters that happen to be in short supply.

Natural limitations
For practical purposes there is a limit on the number of variant families and also a limit on the number of members within any given variant family that can usefully be included in an item vat. Even at the level of the vat, inflating the numbers of items of any given item type by the use of variants might give an unrealistic impression of the robustness of the vat. For this reason, item writers have been given guidelines to follow when planning their item-creation and variant-creation work. The guidelines vary, as appropriate, from one testing program to the other and from one item type to the other.

Next steps

At this time, the writing of variants is considered an integral part of item creation at Educational Testing Service. Not only are variants written in increasing numbers, but also items are created at the outset with an eye toward more efficient variant creation. Moreover, variants are now seen as useful rapid item-creation tools in programs that use paper-based testing exclusively, for they provide the same type of security there as they do for computer-based testing pools. Programs that have a relatively small number of operational paper forms can expand their list of active forms and increase security by making variant test forms--tests composed of items that appear to be on the same topic as items in a parallel test form but with logical or factual structures that lead to different correct answers. There is, once again, a limit to the number of variant items that could usefully be included in the variant test forms--no more than, say, 30 or 40 percent of the items in a given item repository would be members of variant families. This limitation, and a healthy and regular infusion of new items into the repository, would prevent an unseemly shrinkage of the effective domain to be assessed for a given test.

Logical Reasoning Item and Medium Variant

A society can achieve a fair distribution of resources only under conditions of economic growth. There can be no economic growth unless the society guarantees equality of economic opportunity to all of its citizens. Equality of economic opportunity cannot be guaranteed unless a society's government actively works to bring it about.

If the statements given are true, it can be properly concluded from them that

(A) no government can achieve a fair distribution of resources under conditions of economic growth
(B) all societies that guarantee equality of economic opportunity to all of their members are societies that distribute resources fairly
(C) a society can achieve a fair distribution of resources only if its government actively works to bring about equality of economic opportunity
(D) there can be no economic growth in a society unless that society guarantees a fair distribution of resources
(E) some societies that experience economic growth fail to guarantee equality of opportunity to all of their citizens

Statistics:

Form	Base N	Omit	A	B	C	D	E	M-Tot	Scale	E Delta	Crit
3NGR4 P4	1855	14	41	202	1307*	187	54	13.0	3DGR	10.8	XS050
Test Code	Item #	M-O	M-A	M-B	M-C	M-D	M-E	P-Tot	P+	O Delta	RBis
ANLYT 7A	24	12.0	12.0	12.2	13.5	11.8	11.3	0.97	0.72	10.7	0.23

No society can achieve a fair distribution of resources except under conditions of economic growth. There can be no economic growth unless the society guarantees equality of economic opportunity for all its citizens. Equality of economic opportunity cannot be guaranteed unless a society's government actively works to bring it about, and a government can do so only if some of its citizens are willing to make short-term economic sacrifices.

If the statements given are true, which of the following must be true on the basis of them?

(A) A society cannot achieve a fair distribution of resources if none of its citizens are willing to make short-term economic sacrifices.
(B) Unless a society is experiencing economic growth, its government cannot expect any of its citizens to make even short-term economic sacrifices.
(C) Any society whose government actively works to bring about equality of economic opportunity can achieve a fair distribution of resources.
(D) There can be no economic growth in a society unless that society guarantees a fair distribution of resources.
(E) Some societies that experience economic growth fail to guarantee equality of economic opportunity for all their citizens.

Form	Base N	Omit	A	B	C	D	E	M-Tot	Scale	E Delta	Crit
3RGR4 122	920	6	550*	94	150	74	34	13.0	3DGR	12.3	XS050
Test Code	Item #	M-O	M-A	M-B	M-C	M-D	M-E	P-Tot	P+	O Delta	RBis
ANLYT 4A	23	11.3	14.0	11.6	11.4	12.0	11.4	0.99	0.61	11.9	0.38

Analytical Reasoning Stimulus and Medium Variant

Eight volunteers--Grassi, Hong, Jones, Kahn, Lorca, Mott, Nilan, and Patel--are available to clean up two public beaches--Stony Beach and Turtle Beach. Each volunteer will be assigned to exactly one beach, and each beach will be cleaned by four volunteers. The assignment must conform to the following conditions:

Grassi is not assigned to the same beach as Jones.
Jones is not assigned to the same beach as Mott.
If Hong is assigned to Stony Beach, Kahn is assigned to Turtle Beach.
Nilan is assigned to Turtle Beach.

Eight park rangers--F, G, H, J, K, L, M, and N--will each be assigned to one of two patrols, patrol 1 and patrol 2. Each patrol will have four members, and assignments to the patrols are subject to the following constraints:

F cannot be assigned to the same patrol as G.
J must be assigned to the same patrol as K.
G is assigned to patrol 2 if both K and L are assigned to patrol 1; otherwise, G is assigned to patrol 1.

Analytical Reasoning Item

If Grassi is assigned to the same beach as Hong, which of the following can be true?

(A) Hong and Kahn are both assigned to Turtle Beach.
(B) Kahn and Lorca are both assigned to Turtle Beach.
(C) Kahn and Mott are both assigned to Turtle Beach.
(D) Lorca and Mott are both assigned to Turtle Beach.
(E) Mott and Patel are both assigned to Turtle Beach.

Timothy Habick
The Reasoning Group
Educational Testing Service (32-N)
Princeton, NJ 08541
(609) 683-2986
(215) 379-0367

REPRODUCTION RELEASE

(Specific Document)

I. DOCUMENT IDENTIFICATION:

Title: Item Variants in Computer-Based Tests

Author(s): Timothy HABICK

Corporate Source: Educational Testing Service

Publication Date: April 23, 1999

II. REPRODUCTION RELEASE:

In order to disseminate as widely as possible timely and significant materials of interest to the educational community, documents announced in the monthly abstract journal of the ERIC system, *Resources in Education* (RIE), are usually made available to users in microfiche, reproduced paper copy, and electronic media, and sold through the ERIC Document Reproduction Service (EDRS). Credit is given to the source of each document, and, if reproduction release is granted, one of the following notices is affixed to the document.

If permission is granted to reproduce and disseminate the identified document, please CHECK ONE of the following three options and sign at the bottom of the page.

The sample sticker shown below will be affixed to all Level 1 documents	The sample sticker shown below will be affixed to all Level 2A documents	The sample sticker shown below will be affixed to all Level 2B documents
PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY Sample TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) 1	PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL IN MICROFICHE, AND IN ELECTRONIC MEDIA FOR ERIC COLLECTION SUBSCRIBERS ONLY, HAS BEEN GRANTED BY Sample TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) 2A	PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL IN MICROFICHE ONLY HAS BEEN GRANTED BY Sample TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) 2B
Level 1	Level 2A	Level 2B
☑	☐	☐
Check here for Level 1 release, permitting reproduction and dissemination in microfiche or other ERIC archival media (e.g., electronic) *and* paper copy.	Check here for Level 2A release, permitting reproduction and dissemination in microfiche and in electronic media for ERIC archival collection subscribers only	Check here for Level 2B release, permitting reproduction and dissemination in microfiche only

Documents will be processed as indicated provided reproduction quality permits.
If permission to reproduce is granted, but no box is checked, documents will be processed at Level 1.

I hereby grant to the Educational Resources Information Center (ERIC) nonexclusive permission to reproduce and disseminate this document as indicated above. Reproduction from the ERIC microfiche or electronic media by persons other than ERIC employees and its system contractors requires permission from the copyright holder. Exception is made for non-profit reproduction by libraries and other service agencies to satisfy information needs of educators in response to discrete inquiries.

Sign here, please

Signature: Timothy Habick	Printed Name/Position/Title: TIMOTHY HABICK	
Organization/Address: Educational Testing Service	Telephone: 609-683-2986	FAX: 609-497-8785
	E-Mail Address: thabick@ets.org	Date:

(over)

III. DOCUMENT AVAILABILITY INFORMATION (FROM NON-ERIC SOURCE):

If permission to reproduce is not granted to ERIC, or, if you wish ERIC to cite the availability of the document from another source, please provide the following information regarding the availability of the document. (ERIC will not announce a document unless it is publicly available, and a dependable source can be specified. Contributors should also be aware that ERIC selection criteria are significantly more stringent for documents that cannot be made available through EDRS.)

Publisher/Distributor:
Address:
Price:

IV. REFERRAL OF ERIC TO COPYRIGHT/REPRODUCTION RIGHTS HOLDER:

If the right to grant this reproduction release is held by someone other than the addressee, please provide the appropriate name and address:

Name:
Address:

V. WHERE TO SEND THIS FORM:

Send this form to the following ERIC Clearinghouse:

THE UNIVERSITY OF MARYLAND
ERIC CLEARINGHOUSE ON ASSESSMENT AND EVALUATION
1129 SHRIVER LAB, CAMPUS DRIVE
COLLEGE PARK, MD 20742-5701
Attn: Acquisitions

However, if solicited by the ERIC Facility, or if making an unsolicited contribution to ERIC, return this form (and the document being contributed) to:

ERIC Processing and Reference Facility
1100 West Street, 2nd Floor
Laurel, Maryland 20707-3598

Telephone: 301-497-4080
Toll Free: 800-799-3742
FAX: 301-953-0263
e-mail: ericfac@inet.ed.gov
WWW: http://ericfac.piccard.csc.com